Scalability Implications of Virtual Private Networks

Jeremy De Clercq and Olivier Paridaens, Alcatel

ABSTRACT

This article gives an overview of the most promising technologies for service providers to offer virtual private network services. The focus of this article is on the analysis of the scalability implications of these virtual private network mechanisms on existing service provider backbone networks. Very often, when deploying VPN services, service providers will be confronted with a trade-off between scalability and security. VPNs that require site-to-site interconnectivity without strong (cryptographic) security can be deployed in a scalable way based on the network-based VPN model, as long as the interaction between the customer and provider routing dynamics are controlled. VPNs that require strong (end-to-end) cryptographic security should be deployed according to the CPE-based VPN model, using the available IPsec protocol suite.

INTRODUCTION

Virtual private networks (VPNs) have existed for a long time, but until now have been either very expensive or limited in functionality and intensive to manage.

Cost reduction (connectionless instead of connection-oriented networking), new technologies (e.g., multiprotocol label switching, MPLS) and more powerful provider edge devices (e.g., capable of context separation and virtual routing) have enabled service providers (SPs) to build VPNs for their customers in an efficient way. These new cost-effective VPN services are actually often seen as the largest (future) profit generators for IP (converged data and voice) networks.

Although these VPN services seem very promising and have benefited from very optimistic market forecasts, many specialists have serious concerns with regards to the scalability and security of these VPN techniques. This article will give a short overview of the different VPN models that are considered today, and will discuss them in terms of scalability and security.

The deployment of VPNs affects the scalability of the provider networks in terms of memory consumption (amount of code, number of routes to maintain, etc.), processing power (signaling tunnel establishment, updating routing information, etc.) and configuration and management load (upon VPN topology changes, etc.).

As such, a solution is considered "more scalable" than an alternative solution when, for the deployment of the same set of VPNs, fewer devices are affected, memory consumption is smaller, the requested additional processing power is more restricted (resulting in a higher performance), and/or the configuration and management load is smaller.

The question with which SPs are confronted is whether a VPN solution in a specific environment will enable the SP to support enough satisfied VPN customers to justify the investment.

Next to an analysis of VPN solutions on the basis of scalability considerations, the security properties of VPN solutions are analyzed. It shows clearly that strong security has an important cost, and that a trade-off should be made between security and scalability.

SCOPE

VPNs are networks that are perceived as being private networks by the customers using them, but are built over a shared infrastructure owned by an SP. The shared infrastructure consists of the shared backbone and the provider edge devices (PEs). A VPN typically consists of a number of geographically dislocated (private) customer sites that are attached to PEs through customer edge (CE) devices and communicate with each other using a shared backbone (Fig. 1).

THE OLD DAYS

Traditional VPNs exist in two flavors: leased line VPNs and customer-premises-based secure VPNs.

With *traditional leased line VPNs*, the customer sites are interconnected via static (permanent) virtual channels such as asynchronous transfer mode (ATM) or frame relay private virtual connections (PVCs) through a layer 2 backbone network.

The individual sites are connected to the edges of the SP network, and the SP establishes the

0163-6804/02/$17.00 © 2002 IEEE

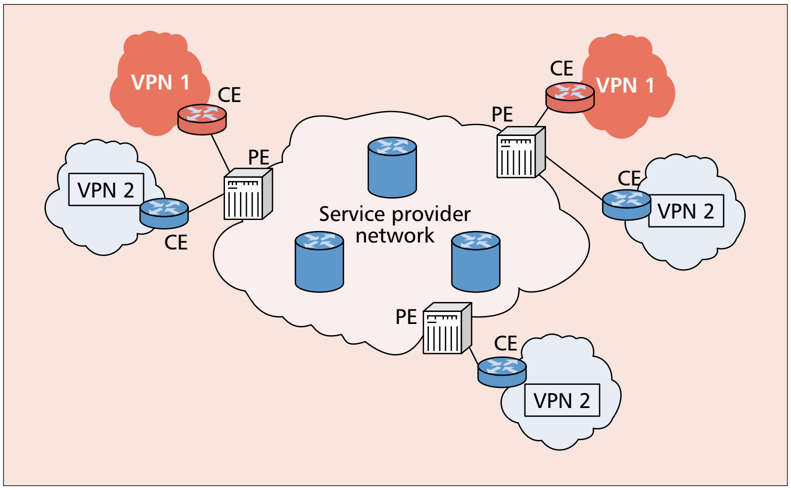

Figure 1. *Reference architecture.*

necessary layer 2 connections. This is a very expensive architecture, in terms of both provisioning as well as configuration and management.

The establishment of a leased-line VPN typically takes a long time, and requires a lot of manpower.

With *customer premises equipment-based* (CPE-based) VPNs, all the VPN functions are implemented at the customer premises. The SP's infrastructure is not involved with any particular VPN function: the routers in the SP's network do not treat VPN IP packets differently from, say, Internet access IP packets.

Customers can buy and deploy dedicated VPN equipment or import software engines on existing routers, gateways, or even personal computers. Since different VPN sites are typically interconnected through the Internet, an unknown and distrusted interconnection of networks, CPE-based VPNs often make use of cryptographic security to protect their intersite traffic.

A drawback of CPE-based VPNs is that they require customers to acquire, configure, and maintain expensive VPN gateways. This implies the presence of highly qualified IT staff.

The New Wave

In the last couple of years, different equipment vendors have proposed a new type of VPNs, and since then, these *network-based IP* VPNs have gained important market interest and an increasing market share. Network-based IP VPNs enable SPs with an IP backbone to offer VPN services to a large number of customers over the same backbone, in a scalable and manageable way, without affecting the existing customer networks. These VPN connectivity services are then often offered in combination with other IP services (e.g., Internet access, firewalls, and IP quality of service, QoS).

With network-based IP VPNs, the customer's routers need not implement VPN-specific functions such as tunneling. Customer sites are connected to PEs that are IP routers (Fig. 1). These PEs need to maintain separate (IP) contexts (with separate IP routing and forwarding tables) for every supported VPN, ensure the distribution of the IP reachability information between distant sites belonging to the same VPN, and intelligently forward the VPN traffic.

The creation and consistent maintenance of separate contexts for different VPNs (as depicted for PE 1 in Fig. 2) makes it possible that:

- Traffic from one specific VPN will not be injected into sites belonging to a different VPN.
- Customer sites can use private addressing, independent of each other and of the SP's addressing realm

Two complementary techniques are used to accomplish this separation of contexts in a scalable way. The first technique is the implementation of virtual routers (VRs) in a PE router, the second the establishment of backbone tunnels between these PE routers.

A VR is a routing and forwarding instance that operates independent of the other VRs in the same PE and of the PE's global routing and

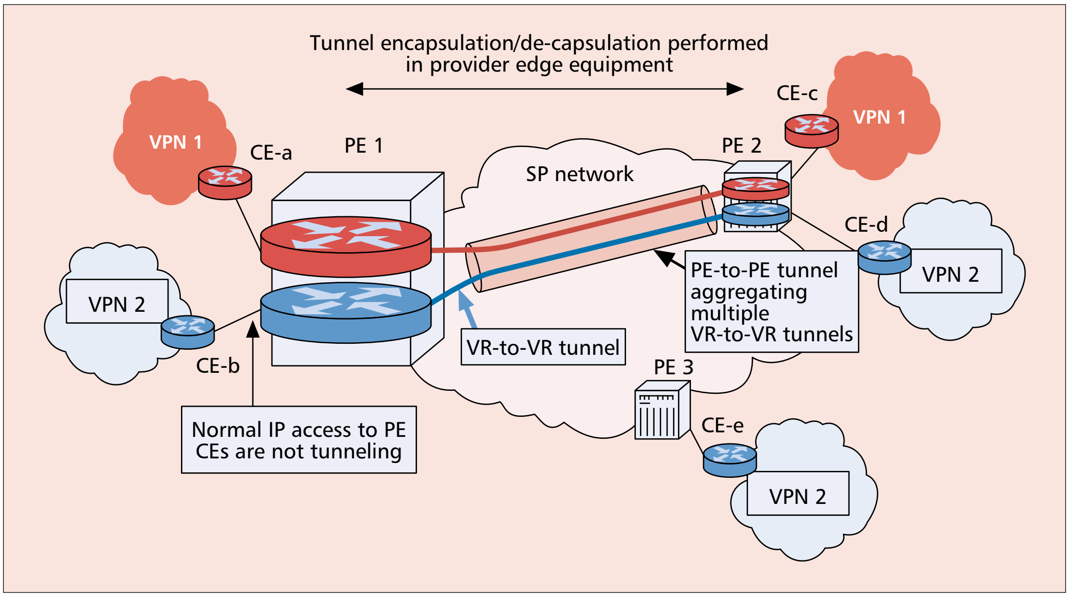

PE-to-PE tunnel aggregating multiple VR-to-VR tunnels

Normal IP access to PE CEs are not tunneling

Figure 2. *Network-based layer 3 IP VPNs.*

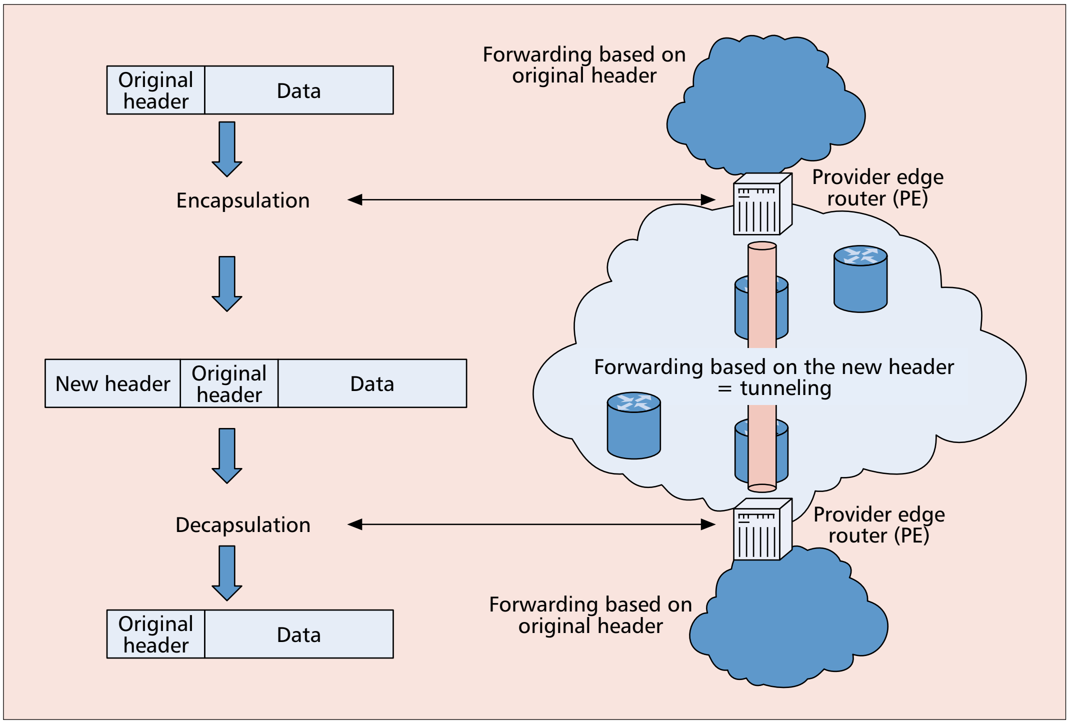
Original header
Data
Encapsulation
New header
Original header
Data
Decapsulation
Original header
Data
Forwarding based on original header
Provider edge router (PE)
Forwarding based on the new header = tunneling
Provider edge router (PE)
Forwarding based on original header

Figure 3. *Tunneling.*

Multiprotocol label switching is a very promising tunneling technique for this application: it offers excellent multiplexing capabilities, a large amount of automation, and very good traffic engineering properties.

forwarding instance. To achieve the separation of VPN contexts, a VR per supported VPN will be deployed in every PE.

The next step is the deployment of backbone tunnels between PEs (Fig. 3). As such, the backbone routers will not need to process the (inner) IP headers of the VPN data units, or maintain separate contexts for the different VPNs.

MPLS is a very promising tunneling technique for this application: it offers excellent multiplexing capabilities, a large amount of automation, and very good traffic engineering properties.

Next to MPLS, other IP tunneling mechanisms such as IP-in-IP tunneling, Layer 2 Tunneling Protocol (L2TP), Generic Route Encapsulation (GRE) tunneling, or IPsec tunneling can be used to tunnel VPN traffic across an SP backbone network.

Other important subjects such as the applicability of the described schemes for IPv6 networks and the transition to these IPv6 networks will not be considered in this article.

Scalability Analysis

Problem Definition

The implementation of VPN-specific functions in the routers of a network will have scalability implications regarding *memory* (per-VPN or per-site state to maintain, especially the VPN IP routes in the routing and forwarding tables), *processing power*, and *management load*.

The general strategy used in all the proposed VPN models is to concentrate the VPN intelligence at the edges of the core network. Tunneling is the mechanism used to accomplish this goal. While, as a result of this tunneling, the core network elements will not be affected by the VPN service, the CE devices or PE routers will need to be extended with several VPN functions.

Let us first come to a simple way of qualifying the scalability of a VPN solution in terms of memory, processing, and management. Important quantifiers are the *number of VPNs* a SP's backbone network can support, the maximal *number of customer sites* one VPN contains, the *maximal amount of information to maintain per customer site*, and the amount of *information to exchange per VPN* (and the frequency of this information exchange).

In terms of scalability of the management system, we discuss:

- The number of concurrent management sessions to maintain
- The effort necessary to add or delete a VPN to/from the SP network or a VPN site to/from a particular established VPN

As an example we will refer to an SP with a backbone network consisting of 50 PEs and supporting 10,000 VPNs. Since tunneling will be used between the edges of the backbone network, the number of hops between PEs is not important. Assume further that a PE can support 500 VPNs, and that a VPN typically consists of 100 sites.

Provider-Provisioned CPE-Based VPNs

Provider-provisioned CPE-based VPNs [1] are CPE-based VPNs where the SP configures and manages the VPN equipment located at the customer premises on behalf of its customers. To be able to configure a large number of VPNs in a scalable way and in an acceptable timeframe, SPs do not manually configure the CEs, but rather use a central management system. A very reliable and secure management system is required for this operation. Traffic protection is achieved by establishing secure customer-to-customer IPsec tunnels (Fig. 4). For a fully meshed VPN, a different IPsec tunnel must be established between each pair of CEs belonging to the same VPN.

A very reliable and secure management system is required for this operation. Traffic protection is achieved by establishing secure customer-to-customer IPsec tunnels. For a fully meshed VPN, a different IPsec tunnel must be established between each pair of CEs belonging to the same VPN.

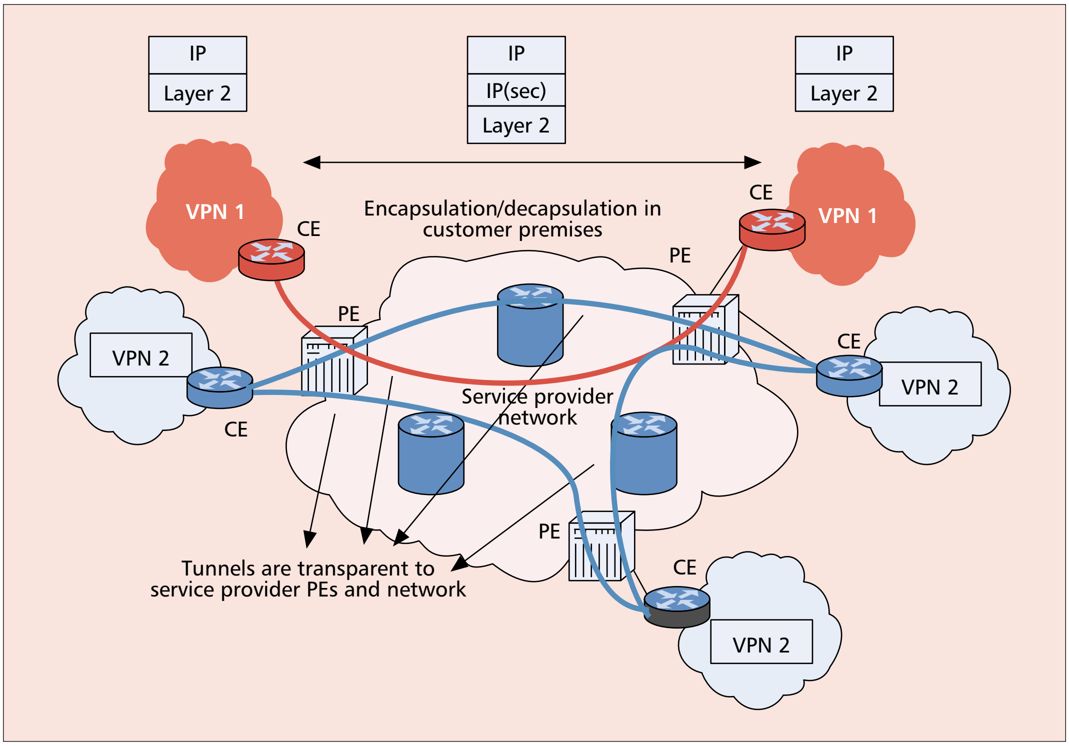
IP
Layer 2
IP
IP(sec)
Layer 2
IP
Layer 2
VPN 1
CE
Encapsulation/decapsulation in customer premises
CE
VPN 1
PE
PE
VPN 2
CE
CE
VPN 2
Service provider network
PE
CE
Tunnels are transparent to service provider PEs and network
VPN 2

Figure 4. *CPE-based IP VPNs.*

In terms of scalability:

- The management system has to scale to support a very large number of VPN customers and a very large number of sites per VPN customer.
- Every CE (in a meshed VPN topology) or alternatively one VPN gateway at the customer premises headquarters (in a star topology) needs to be able to support a large number of (preferably encrypted) connections in terms of data processing and connection maintenance.

Regarding the example configuration, this would mean:

- The SP's management system needs to be able to concurrently support up to 1 million management sessions.
- The CE of a hub site needs to be able to concurrently support 99 VPN connections (tunnels).

Thus, it is clear that provider provisioned CPE-based VPNs put a large load on the SP's management system.

Network-Based Layer 3 VPNs

The number of network-based layer 3 VPNs a network can support is a relative number since it is highly dependent on the number of sites belonging to the VPN and the amount of information to maintain per VPN site.

Since not all the PEs will have to support all the VPNs the provider is supporting (Fig. 1), it is more interesting to focus the scalability discussion on the PE (the edge router of the backbone network), and to take the following parameters into account: the number of VPN contexts supported per PE, the amount of state to maintain per VPN context in a PE, and the amount of control information a PE has to process and exchange per VPN.

Per-VPN Routing State — Implementing IP VPN functionality in PE routers means creating a separate context for every VPN the particular PE will support. This separate context will need to maintain state for that particular VPN, and will need to have dedicated routing and forwarding instances. These separate contexts are often called virtual routers (VRs) or virtual routing and forwarding instances (VRFs). VRs can be implemented in software on an existing platform, or alternatively one can optimize hardware for VR support.

The number of VRs a PE can support is a key parameter for the scalability of a VPN service. These VRs need to maintain routing and forwarding tables, and support routing protocols that run in addition to the PE's global (Internet) routing and forwarding tasks.

The manual configuration of routing information in VRs is an impossible task in a dynamic environment where customer routes change, VPN sites are added and deleted, and so on, especially for a large SP network supporting a large number of VPNs. Therefore, it is advisable to involve the VRs in the process of exchanging routing information with the customer's network.

When routing protocol interaction is supported between the customer's sites and the VRs, care must be taken to control the routing information inserted in a VR. Badly designed customer networks can lead to a huge amount of reachability information to be maintained in all VRs. Not only will these routes be maintained in the directly attached VR, but the information will also be distributed to the VRs from the same VPN in other PEs, and also result in extra VPN state. In many cases, a PE already needs to maintain a (large) global (Internet) routing/forwarding table, leading to possible scalability concerns. When in addition a large number of VRs must be deployed for VPN support, scalability can only be ensured when the routing and forwarding tables can be kept to a reasonable size.

Therefore, customers need to be helped or obliged to intelligently design their private networks, and thus to limit the number of routes introduced in the VRs. When customers properly assign subnets to VPN sites, the reachability information can be summarized or aggregated into a small amount of routes in the corresponding VR.

In terms of the quantitative example to which this article refers, a PE will need to be able to support 500 VRs, and therefore 500 independent routing and forwarding tables, each maintaining a number of routes dependent on the supported VPNs (e.g., if every site announces 10 routes to its VR, a single PE will need to maintain 500,000 VPN routes!).

Distribution of VPN Routing Information — To ensure intersite VPN reachability, the VPN reachability information needs to be distributed over the shared backbone to the other sites of the same VPN. The goal is again not to affect the backbone network elements, and to concentrate the intelligence into the PEs. Two mechanisms are currently used for this.

In the first model (the *peering* model), the VRs in different PEs that belong to the same VPN have a direct peering relationship in terms of reachability information distribution: a routing protocol instance runs between them through the same VPN tunnels used for the data plane [2].

In the second model (the *overlay* model [3]), either a new instance of a core edge-to-edge routing protocol (e.g., Border Gateway Protocol 4, BGP-4) is deployed with its own VPN multiplexing and filtering capabilities, or the VPN routing information is piggybacked on a PE's border routing instance already used for global (Internet) routing information exchange.

Although these models are conceptually different, their impact on the scalability of routing processing and the information exchanges is identical: for the same VPN architecture and topology, the number of routes to process and distribute is identical.

Both the *number of routes* that need to be maintained and the required *processing* involved with the propagation of routing changes are independent of whether an overlay or peering model is used. The former solely depends on the amount of reachability information advertised by the CEs; the latter is dependent on the chosen routing protocol and the frequency of the routing changes advertised by the CEs.

Generally speaking, the scalability of the VPN routing distribution is not dependent on the number of sites per VPN, it only depends on the number of PEs that serve a common VPN.

In terms of the example, a PE needs to support less than 24,500 (49 × 500) routing update sessions.

The *frequency* of the routing information processing and distribution has a major impact on the scalability of the system. An intelligent implementation will make sure that the frequency of the information distribution to the peering edge routers will be minimized.

Two factors have an impact on this:

- The routing protocol instance that distributes the information between different PEs
- The routing protocol interaction between the PE and the customer sites

The former is largely dependent on the chosen routing protocol, while the latter can be controlled locally.

Indeed, we pointed out earlier that the dynamics of the customer's routing information exchange can have an impact on the amount of state to maintain, but the frequency of information distribution and stability of the customer routing domain can have a similar impact on the scalability and stability of the SP network.

The VRs need to be designed such that possible instabilities such as routing update message storms and loops in routing advertisements in the customer domain are detected and not propagated through the provider's network.

Also, very dynamic environments like dial-in services can affect the scalability of the provider network. Indeed, in theory every new customer that dials in to a certain VPN receives a new IP address that needs to be propagated within the whole existing VPN (to all impacted PEs, over the backbone network). Since the dial-in rate can be quite high for certain applications, this is a very unscalable solution. In this particular situation, the use of a pool of IP addresses per PE is recommended. This pool can then be distributed among the other PEs as an address prefix, even before the first customer dials in. As such, the event of a user dialing in has only a local impact on the directly attached PE.

Management Complexity — Another scalability concern introduced with VPN services is whether, when a large number of VPNs is supported, the configuration of a new VPN or the addition of a new site to an existing VPN, is a manageable task.

In practice, every PE needs to have a global view of the topology of every supported VPN: a PE needs to know, for every supported VPN, what other PEs serve the same VPN. Theoretically, this information needs to be configured in every PE of an SP's network: when a new site belonging to a certain VPN needs to be added to a particular PE, every other PE can be manually configured with this information.

For large SP networks with a large number of PEs, this is an intensive and error-prone task leading to long response times. Therefore, most VPN models use an "auto-discovery" mechanism that takes care of this issue. This can be achieved via an intelligent automatic management system, or a protocol that runs directly between the PEs. Two protocols being proposed for this task are BGP-4 and the Domain Name Server (DNS) protocol.

As such, the addition of a new VPN site only requires configuration of the PE to which the new site is attached. The other PEs will auto-discover this new information.

It is not believed that the use of auto-discovery schemes will impact the scalability of VPN solutions, since the addition/deletion of VPN sites is not expected to be very frequent.

To assure inter-site VPN reachability, the VPN reachability information needs to be distributed over the shared backbone to the other sites of the same VPN. The goal is again not to affect the backbone network elements, and to concentrate the intelligence into the PEs.

The most important security requirement is the protection of the customer's traffic over the SP's backbone. Indeed, the SP network is not (necessarily) fully trustable since it may rely on other providers' networks to build the customer's VPN.

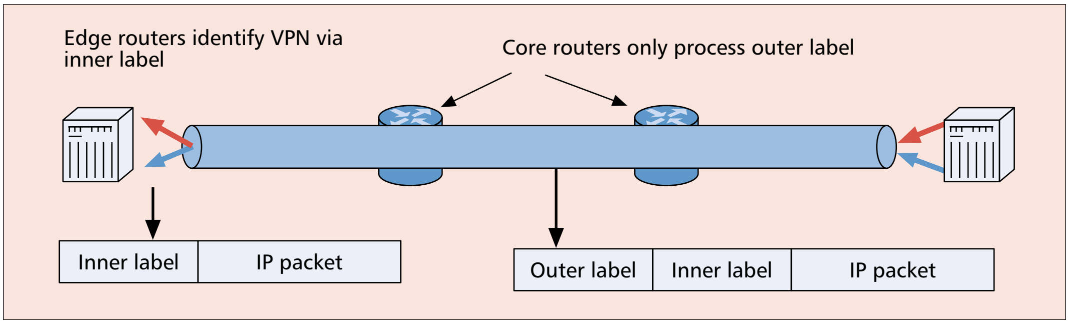
Edge routers identify VPN via inner label
Core routers only process outer label
Inner label
IP packet
Outer label
Inner label
IP packet

Figure 5. *Multiplexing of MPLS tunnels.*

Maintaining VPN Tunnels — Next to the necessary support of multiple VPN contexts and routing instances, and the implication for the network management system, another VPN feature that might have implications on the scalability of the PE router is the fact that tunnels need to be established and maintained with the other PE routers. Conceptually, a tunnel between two provider edge routers is needed for every VPN context they have in common. Now, in connectionless networking, state for every tunnel is maintained only at the edges. This is the case, for example, in IPsec, IP-in-IP, GRE, and L2TP.

MPLS, on the other hand, as a connection-oriented technology, introduces state in every node that supports a certain connection. This could lead to serious scalability problems in core networks where large numbers of MPLS tunnels are requested. The multiplexing of tunnels into larger tunnel pipes deals with this issue. VPN architectures that make use of the MPLS technology in the backbone network are good examples of the multiplexing of tunnels. The MPLS tunnels that need to be established between the VRs of two specific PEs are then multiplexed into a large MPLS pipe between these PEs. As such, the backbone routers need only maintain state for these (few) large MPLS pipes (Fig. 5).

It is only the PE routers that need to maintain tunnel information for the inter-VR tunnels. Note that this information is very limited, since no specific processing like encryption is required.

In fact, the information a PE needs to maintain is one MPLS label per deployed VR for incoming MPLS packets (the *inner label* in Fig. 5), one label per peer PE (the *outer label* in Fig. 5) for the pipe label switched path (LSP) to use, and one MPLS label per common VPN per peer PE for outgoing MPLS packets (the *inner label* in Fig. 5).

In terms of the example, a PE would need to maintain 500 *incoming labels*, and less than 24,549 (49 + 500 × 49) *outgoing labels*. A provider core router, on the other hand, would only need to support up to 2450 labels, independent of the number of supported VPNs.

Security Analysis

Security Requirements and Implications

Because there is often a trade-off between security and scalability, we here list some important security features that must be considered when comparing different types of VPN solutions.

The most important security requirement is the protection (confidentiality, integrity, authentication, replay detection) of the customer's traffic over the SP's backbone. Indeed, the SP network is not (necessarily) fully trustable since it may rely on other providers' networks to build the customer's VPN. Even if the SP owns the complete network, the customer may protect itself against any eavesdropping on the SP network. Encrypting the customer's VPN traffic prevents other customers in other VPNs from accessing the data in case of (accidental or deliberate) misrouting within the SP network (in Fig. 1, it should not be possible for VPN1 users to read VPN2 traffic even if the traffic was accidentally misrouted). Next to encryption, authentication of VPN traffic is also useful to filter incoming traffic (at the CE or PE) to ensure authenticity of the received traffic. The classical technique adopted to build secure VPNs is to make use of IPsec to secure the IP or MPLS traffic.

In order to achieve maximum security, it is recommended to use end-to-end encryption. For VPNs, there are basically two security models: a network-based (i.e., PE-to-PE) model and a customer-based (i.e., CE-to-CE) security model. CE-to-CE security provides a secure VPN that is closer to the end-to-end model, and hence is usually the recommended solution from a security point of view.

The PE-to-PE security model does not inherently cover the PE-CE leg, although this link can be separately secured at the cost of extra processing in the PE and VPN traffic being "accessible" within the PE. This model also fits well when providing additional services such as network address translation, firewall, and so on in the PE under control of the SP.

When the CE-PE link relies on a shared infrastructure, securing that leg can be more of a concern. It is interesting to note that although the CE-to-CE secure VPN model protects the traffic over this path, it does not fully protect against VPN resource spoofing by invalid users. An intruder can still inject traffic on the CE-PE link, which will be forwarded by the PE since it is unable to authenticate traffic received from the CE (authentication of VPN traffic is done at the CE level).

Apart from protecting the customer's VPN traffic, it is also important to consider protection of the SP infrastructure and overall VPN management.

In terms of scalability, it is difficult to compare both security models, since different network elements are affected.

For a VPN consisting of 100 sites, the CE-to-CE secure VPN model basically requires the establishment of 99 secure channels in each CE.

This number can be reduced to a single secure channel per CE in a star VPN topology (with the central site maintaining 99 secure channels). Regarding the overall scalability issue, the impact on the performance can be high for low-range CE devices that must manage a large number of different security channels and process the data traffic accordingly.

Note that the application of encryption security for VPNs is independent of whether the VPN tunneling is IP- or MPLS-based.

Cost of Security

Applying cryptographic security mechanisms in VPN solutions implies an increase of memory consumption (state that has to be maintained for every supported security association, etc.), and a serious increase in processing (cryptographic functions are very demanding in terms of processing, secure channels need to be frequently renegotiated, etc.).

This results in a trade-off between security and scalability: for the same amount of memory and processing power in a specific device, an increase of the number of cryptographically secured VPN connections results in a decrease in the total number of VPN connections or in performance.

A Security Model for Network-Based Layer 3 VPNs

For network-based layer 3 VPNs, the customer can still choose to secure VPN traffic on a CE-to-CE basis, independent of the SP. Another method is to rely on the SP securing the VPN traffic between PEs using IPsec.

When secure channels are established on a PE-to-PE basis, two techniques are possible:

- Establishing secure channels on a per-VPN basis; IPsec in tunnel mode can indeed be used to establish security channels between VRs. This has serious scalability implications, since it is necessary to signal and maintain multiple secure channels (IPsec security associations) between each pair of PEs (24,500/PE in the example).
- Establishing secure channels on a per-remote-PE basis; this solution is much more scalable, since the number of security associations to maintain per PE does not depend on the number of VPNs, but only on the number of PEs (example: 49/PE)

Since both discussed techniques offer the same protection (in both cases the end nodes, the PEs, are under control of the same entity, the SP), it is sensible to design solutions that enable the use of security channels on a per-remote-PE basis.

The recommended solution consists of applying IPsec in transport mode to IP-tunneled MPLS packets that carry the actual customer IP traffic [4]. The MPLS labels are then used to differentiate the traffic from different VPNs secured with the same security channel. As such, the number of secure channels to establish and manage is greatly reduced.

Another solution than this MPLS-in-IP/IPsec would be to secure the MPLS data plane by providing security mechanisms in MPLS itself. There would be no difference, though, as to the number of security associations to maintain.

Conclusion

This article gives an overview of some of the most promising IP VPN mechanisms and discusses them in regard to *scalability* and *security*.

For security-sensitive applications, it is recommended to apply cryptographic security on an end-to-end basis (e.g., in a CE-based model).

Network-based IP VPNs offer the possibility to deploy scalable VPN services in large SP networks. It is recommended, however, to control and limit the routing information exchange between the SP core network and the customer networks (and the rate of this information exchange).

In addition, when it is necessary to apply cryptographic security on a network level, it is recommended to limit the number of signaled and maintained security associations by deploying secure channels on a per-remote PE basis, instead of on a per-VPN basis, according to one of the models this article proposes.

Network-based IP VPNs offer the possibility to deploy scalable VPN services in large SP networks. It is recommended though to control and limit the routing information exchange between the SP core network and the customer networks.

References

[1] J. De Clercq *et al.*, "A Framework for Provider-Provisioned CPE-Based VPNs Using IPsec," IETF draft, http://search.ietf.org:80/html.charters/ppvpn-charter.html July 2001, Feb. 2002, work in progress.

[2] H. Ould-Brahim *et al.*, "Network based IP VPN Architecture using Virtual Routers," IETF draft, http://search.ietf.org:80/html.charters/ppvpn-charter.html, July 2001, Feb. 2002, work in progress.

[3] E. Rosen *et al.*, "BGP/MPLS VPNs," RFC 2547, July 2001.

[4] E. Rosen *et al.*, "Use of PE-PE IPsec in 2547 VPNs," IETF-draft, http://search.ietf.org:80/html.charters/ppvpn-charter.html, July 2001, Feb. 2002, work in progress.

Biographies

Jeremy De Clercq (jeremy.de_clercq@alcatel.be) graduated in electro-technical engineering, specialized in communication technologies, at the University of Ghent in 1999. He started his professional carrier in the Corporate Research Center of Alcatel. He currently works in Alcatel's Network Strategy Group, providing internal strategy and consultancy on network architectures. He is also active in standardization, mainly in the Internet Engineering Task Force (IETF).

Olivier Paridaens (olivier.paridaens@alcatel.be) graduated in computer sciences from Université Libre de Bruxelles in 1989. He worked for over 10 years in a university department doing research and consultancy on networking technologies and ICT security. Since then he joined Alcatel where he leads a team providing internal strategy and consultancy on security technologies. His main current activities are related to security in NGN, UMTS, and VPN environments.